



Europäisches Patentamt

(19)

European Patent Office

Office européen des brevets



(11)

EP 0 910 047 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:

21.04.1999 Bulletin 1999/16

(51) Int. Cl.<sup>6</sup>: G06T 15/10

(21) Application number: 98307790.0

(22) Date of filing: 25.09.1998

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 15.10.1997 US 958129

(71) Applicant:

DIGITAL EQUIPMENT CORPORATION  
Maynard, Massachusetts 01754 (US)

(72) Inventors:

- Jouppi, Norman, P  
Palo Alto, California 94306 (US)
- McCormack, Joel, M  
Boulder, Colorado 80302 (US)

(74) Representative:

Charig, Raymond Julian et al  
Eric Potter Clarkson,  
Park View House,  
58 The Ropewalk  
Nottingham NG1 5DD (GB)

## (54) Full-scene antialiasing using improved supersampling techniques

(57) A method and an apparatus reduces aliasing artifacts in images defined by pixels. A pixel is partitioned into subpixel locations from which sample points are selected. A fragment of the image is determined to be visible at least one of the sample points. A fragment value associated with that fragment is stored. Each sample point at which the fragment is visible is linked to the stored fragment value. A color of the pixel is computed from the stored fragment values to reduce the aliasing artifacts in the image.




FIG. 1

EP 0 910 047 A2

1

EP 0 910 047 A2

2

**Description****Field of the Invention**

[0001] This invention relates generally to computer graphics, and more particularly to a method and apparatus for reducing aliasing artifacts in images defined by pixels.

**Background**

[0002] In pixel-based display systems, visual artifacts such as jagged lines are due to an effect known as aliasing. Aliasing occurs because of the discrete nature of pixels, the smallest picture element addressable on a display screen. Pixels are arranged on a display screen as an rectangular array of points. Aliasing artifacts can occur when an entire pixel is given a light intensity or color based upon an insufficient sample of points within that pixel. Techniques to minimize aliasing artifacts are referred to as antialiasing.

[0003] One antialiasing technique is supersampling. Supersampling involves taking more samples of an image than there are pixels to be displayed. Such samples are taken at subpixel positions within each pixel. The color and intensity displayed for each pixel comes from combining the subpixel samples.

[0004] The quality of the antialiasing obtained by supersampling is affected by the degree to which pixels are divided into subpixels, and the number of subpixels that are sampled for each pixel. Generally, the more subpixels per pixel, the finer is the resolution of the image; and the more subpixels per pixel sampled, the better is the antialiasing effect. An 8 x 8 grid of subpixels offers four times as many possible sample locations than a 4 x 4 grid of subpixels. Also, sampling all sixty-four subpixels of an 8 x 8 grid provides sixty-four times more pixel data than sampling only one of the sixty-four subpixel sites.

[0005] However, sampling at every subpixel position, referred to as full scene supersampling, is costly to implement. Each subpixel sample requires memory, and consumes memory bandwidth. Also, the amount of computation needed to derive the displayed pixel value increases with the number of samples taken. Accordingly, the problems of full scene supersampling have led to sparse supersampling in which only a portion of the subpixel positions are sampled. In the 8 x 8 grid, for example, only eight of the sixty-four subpixels might be sampled to produce a display value for a pixel.

[0006] Eight subpixel samples per pixel still take up considerable memory. If each pixel needs eight bytes per sample, then eight samples need sixty-four bytes. For a 1600 x 1200 pixel display, this would require 123 Mbytes of memory.

[0007] Thus, there is a need for a method and an apparatus that reduce the memory capacity and bandwidth requirements associated with prior art antialiasing

techniques while effectively reducing aliasing artifacts in images.

**Summary of the Invention**

[0008] This invention is for the purpose of minimizing aliasing artifacts in graphic images using improved supersampling techniques. The invention enables to effectively minimize these artifacts while reducing the memory capacity and bandwidth requirements associated with prior art antialiasing techniques.

[0009] The present invention, in its broad form, resides in a method, a memory, and an apparatus as in claims 1, 11, and 14 respectively for reducing aliasing artifacts in an image defined by pixels. As described herein, the method, selects subpixel locations in a pixel as sample points. A fragment value associated with a fragment of the image is stored. This fragment covers one or more of the selected sample points. Each covered sample point is linked to the stored fragment value to enable the generation of a color of the pixel using the fragment value. The generated color improves the perceived quality of the image by reducing aliasing artifacts.

[0010] Preferably, an index is associated with each covered sample point, and a value stored in each index points to the stored fragment value.

[0011] Further, a bit pattern is associated with the stored fragment value. Each bit in the pattern is associated with one of the sample points. A value is stored in each bit, indicating whether or not the sample point associated with that bit is covered by the fragment.

[0012] As described hereinafter, a new visible fragment is processed, and the stored fragment value is replaced by the fragment value associated with the new fragment.

[0013] Advantageously, addresses of memory are allocated to the pixel for storing a predetermined number of fragment values, the predetermined number being less than the number of sample points.

[0014] In terms of the apparatus, the arrangement comprises a graphics device that selects subpixel positions in a pixel as sample points, and a memory coupled to the graphics device. The memory stores a fragment value associated with a fragment of the image. This fragment covers one or more of the sample points in the pixel. The graphics device links each covered sample point to the stored fragment value to enable the generation of a color of the pixel using the fragment value. When the pixel is rendered, the generated color improves the perceived quality of the image by reducing aliasing artifacts.

[0015] In terms of a memory for storing data used to compose a color for a pixel that is partitioned into subpixels, the described arrangement comprises a plurality of addresses that each store a fragment value. Each stored fragment value includes a color value and is associated with a fragment of an image that covers one

of the subpixels. Another address stores a plurality of indices. Each index is associated with one of the subpixels and includes an index value that points to one of the plurality of addresses storing a particular fragment value to link the subpixel associated with that index with the color value of the particular fragment value.

### Brief Description of the Drawings

[0016] A more detailed understanding of the invention may be had from the following description of a preferred embodiment, given by way of example, and to be understood with reference to the accompanying drawing wherein:

- FIG. 1 is a block diagram of an exemplary computer graphics system that can be used to practice the invention;
- Figs. 2A-2C represent various subdivisions of a pixel into subpixels, and illustrate exemplary sparse supersampling patterns that can be used to sample the subpixels;
- FIG. 3 represents an exemplary linking of subpixel samples in one of the supersampling patterns of FIG. 2A-2C to two fragment triples stored in a pixel memory;
- FIG. 4 represents another linking of subpixel samples for when a third fragment appears in a pixel;
- FIG. 5A-5C illustrate alternative linkings of subpixels samples for when a third fragment appears in a pixel;
- Figs. 6A-6B illustrate a logical representation of the pixel memory including indices to the stored fragment triples;
- Figs. 6C-6D illustrate a logical representation of the pixel memory including coverage masks associated with the stored fragment triples; and
- Fig. 7 illustrates a flow diagram describing an exemplary process using the present invention.

### DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

#### System Overview

[0017] FIG. 1 shows a computer system 100 that can generate monochrome or multicolor 2-dimensional (2D) and 3-dimensional (3D) graphic images for display according to the principles of the present invention. The computer system 100 can be one of a variety of raster graphics systems including, for example, a personal computer, a workstation, or a mainframe.

[0018] In the computer system 100, a system chipset 104 provides an interface among a processing unit 102, a main memory 106, a graphics accelerator 108 and devices (not shown) on an I/O bus 110. The processing unit 102 is coupled to the system chipset 104 by the host bus 112 and includes a central processing unit

(CPU) 118. The main memory 106 interfaces to the system chipset 104 by bus 114.

[0019] The graphics accelerator 108 is coupled to the system chipset 104 by a bus 116, by which the graphics accelerator 108 can receive graphics commands to render graphical images. A graphics memory 122 and a display device 126 are coupled to the graphics accelerator 108; the graphics memory 122 is coupled by bus 124, and the display device 126, by bus 127. The display device 128 includes a cathode ray tube (CRT) raster display monitor 128 with a display surface or screen 130. The CRT 128 produces color images; but the invention can also be practiced with a monochrome monitor to display gray-scale images or with a printer that prints black and white or color images.

[0020] An image 132 appears on the display screen 130 by illuminating a particular pattern of individual points called pixels 134. The image 132, for example, can be 2D alphanumeric characters or a 3D scene filled with objects. The display screen 130 includes a two-dimensional array of such pixels 134. The array size of display screens 130 can vary widely. Examples of display screen 130 sizes include 1024 x 768 and 1320 x 1200 pixels. For the purposes of practicing the invention, the display device 126 may be any other pixel-based display such as a liquid-crystal display or a dot matrix printer.

[0021] The graphics memory 122 includes storage elements for storing an encoded version of the graphical image 132. There is a direct correspondence between the storage elements and each pixel 134 on the display screen 130. The storage elements are allocated to store data representing each pixel 134, hereafter referred to as pixel data. For example, five bytes may be used to encode a color representation for each pixel.

[0022] The values stored in the storage elements for a particular pixel controls the color of the particular pixel 134 on the screen 130. By "color", it is to be understood that the brightness or intensity of the pixel 134 is also intended. Pixel data can translate directly into colors or into indices to access a color lookup table.

[0023] During operation, the computer system 100 can issue graphics commands that request an object to be displayed. The graphics accelerator 108 executes the graphics commands, converting the object into primitives and then into fragments. A primitive is a graphical structure, such as a line, a triangle, a circle, or a surface patch of a solid shape, which can be used to build more complex structures. A fragment is a 2D polygon created by clipping a primitive of the image 132, such as a line, triangle, or circle, to the boundaries of the pixel 134. A more detailed description of fragments is provided by Loren Carpenter in "The A-buffer, an Antialiased Hidden Surface Method", Computer Graphics Vol. 18, No. 3, 1984, pp. 103-107, incorporated by reference herein. There, techniques merge fragments into a fragment list when the fragments are from the same object or surface of the image. Here, the frag-

5

EP 0 910 047 A2

6

ments that are combined to produce the color of a pixel can have a different relationship to each other: that is, the fragments can be from different objects or surfaces of the image 132.

[0024] The graphics accelerator 108 renders the fragments, and loads the pixel data corresponding to the fragments into the appropriate storage elements of the graphics memory 122. The pixel data can be transferred into the graphics memory 122 from the main memory 106 via busses 112, 114, 116, and 124, or written directly into the graphics memory 122 by the graphics accelerator 108.

[0025] To display the image 132, the CRT monitor 128 projects a beam onto the screen 130. In synchrony, the pixel data are read out of the graphics memory 122 as the beam scans the screen 130. The CRT monitor 128 renders the pixel data as illuminated points of color on the display screen 130.

[0026] Figs. 2A-2C illustrate various exemplary subdivisions of a pixel 134. FIG. 2A shows pixel 134 divided into a 4 x 4 array 200 of evenly spaced points called subpixels 206; FIG. 2B shows an 8 x 8 array 202 of subpixels 206; and FIG. 2C shows a 16 x 16 array 204. Dividing a pixel 134 into subpixels 206 provides multiple points at which the image 132 covering that pixel 134 can be sampled. For reference, the center 201 of the pixel 134 is shown as an X.

[0027] Generally, the more subpixels 206 there are in the array, the greater the resolution of the pixel 134. Thus, the displayed color of the pixel 134 does not rely entirely on one sample point, but upon several subpixel samples 206. Methods for calculating a pixel value from multiple sample points are well known in the art.

[0028] Known implementations sampled at every subpixel 206 in a pixel 134. While, theoretically, such full scene supersampling presented opportunities for attaining high resolution, the technique unnecessarily consumed memory resources. Each sampled subpixel 206 required memory resources to store and use the sampled data. Thus, fully sampling the 4 x 4 array 200 of subpixels 206 required memory storage for sixteen samples, in addition to the typical memory requirements for each pixel 134. If the sixteen samples each required, for example, eight bytes of storage, then implementing full scene supersampling could require an additional 295 MBytes of memory for a 1920 x 1200 pixel display screen 130. The 16 x 16 array 204, which requires storage for 256 samples, needs sixteen times as much memory.

[0029] Accordingly, recent modern implementations do not sample at every subpixel 206. Rather, those subpixels 206 which are sampled are sparsely distributed in the subpixel array. In general, the antialiasing results were almost as effective for such sparse supersampling as for the full scene supersampling technique.

[0030] Figs. 2A-2C each illustrate an exemplary sparse supersampling pattern 210, 220, 230 that can be used to sample the subpixels 206 of the corresponding

subpixel array. The illustrated exemplary sample patterns 210, 220, 230 each have N samples distributed uniformly throughout an N x N subpixel array with exactly one subpixel sample in any particular row and in any particular column.

[0031] The sampling pattern 210 has four subpixel samples S1-S4 (N equals 4). For sampling pattern 220, N equals 8, and the eight subpixel samples 222 are S1-S8. For sampling pattern 230, N equals 16, and the sixteen subpixel samples 232 are S1-S16. The sampling pattern 210, 220, 230 can be repeated for every pixel 134 on the display screen 130. Various other sampling patterns can be used to practice the principles of the invention.

[0032] Although sparse supersampling uses less memory than full scene supersampling, considerable amounts of additional memory are still required. For example, when N equals 4, a 1920 x 1200 pixel screen 130 still needs eight bytes storage for each of four subpixel samples. This requires an additional 74 Mbytes of pixel data. The requirements are doubled and quadrupled when N equals 8 and 16, respectively.

[0033] The described arrangement can reduce the storage requirements even more than such sparse supersampling without reducing the number of subpixel samples for an N x N subpixel array. The arrangement described herein relies upon the observation that typically only a few fragments of the image 132 are visible within a given pixel.

[0034] For static and animated images, the antialiasing effects achieved by eight sparse supersamples in the 8 x 8 array 202 appear significantly better than for four samples in the 4 x 4 array 200. But differences between sixteen samples in the 16 x 16 array 204 and eight samples in the 8 x 8 array 202 may be indistinguishable.

[0035] FIG. 3 shows an exemplary pixel 300 that is part of the image 132 and is subdivided into a 4 x 4 subpixel array 200. The pixel 300 has four sampling positions according to sampling pattern 210 of FIG. 2A. Two fragments 301, 302 are in pixel 300. Each fragment 301, 302 is associated with a fragment value, called a fragment triple 310, 312. For example, in FIG. 3, fragment triple 310 is associated with fragment 302, and fragment triple 312 with fragment 301.

[0036] Fragment values are called fragment triples because each fragment triple 310, 312 includes three values: a color value 304, a Z-depth value 306, and a stencil value 308. The color value 304 represents the color and opacity of the corresponding fragment. The Z-depth value 306 represents a Z-coordinate value of the corresponding fragment along a Z-axis that is perpendicular to the image 132 to provide 3D depth. The stencil value 308 can be used to group or identify sets of fragments of the image 132, or to logically or arithmetically process or count operations upon fragments, or for other purposes known to those skilled in the art.

[0037] In the preferred embodiment, the exemplary

fragment triples 310, 312 each use five bytes to represent the color 304, three bytes for the Z-depth 306 and one byte for the stencil 308. The five color 304 bytes accommodate four 10-bit color channels: Red, Green, Blue, and Alpha.

[0038] The color of a fragment is expressed by the combination of the values stored in the Red, Green and Blue (RGB) channels. The value stored in each RGB channel indicates the intensity (or brightness) of that color channel. Low values correspond to low intensity, dark colors; high values correspond to high intensity, light colors. Various methods for producing the color by combining the RGB values are well known in the art.

[0039] The opacity of the fragment is expressed by the value stored in the Alpha channel. For example, a 1.0 value (i.e., all 10 Alpha-channel bits are 1) indicates that the associated fragment is opaque, a 0.0 value indicates that the fragment is invisible, i.e., completely transparent, and values between 0.0 and 1.0 indicate degrees of transparency.

[0040] Memory is allocated to each pixel 134 for storing a predetermined number of fragment triples. This memory can be either graphics memory 122, as shown in FIG. 3, or main memory 106. In the example shown in FIG. 3, the pixel memory 314 is allocated for one particular pixel 300. Conceivably, a group of pixels, like a 2 x 2 array of pixels 134, can share a particular pixel memory 314. Any fragment triples stored in the pixel memory 314 would be used by each pixel 134 in the group, rather than by only one particular pixel 300. This can save more memory than storing a predetermined number of fragments for every pixel 134, particularly for portions of the image 132 that change color and Z-depth gradually.

[0041] Alternatively, memory for storing fragment triples can be dynamically allocated to each pixel 134 rather than fixed to a predetermined number. Here, a variable number of fragment triples can be stored for each pixel 134, the graphics accelerator 108 allocating memory to the pixel 134 as needed, presuming there is still available pixel memory in the system 100. Another method combines aspects of both above-described methods, allocating memory to each pixel 134 for storing a predetermined number of fragment triples, and dynamically allocating additional memory to a particular pixel 134 when needed to store a fragment triple beyond the predetermined number.

[0042] The exemplary embodiment shown in FIG. 3 stores two fragment triples 310, 312 in the pixel memory 314. These fragment triples 310, 312 are associated with the fragments 301, 302 that cover the pixel 300. Before the fragments 301, 302 appear in the pixel 300, the pixel memory 314 can be initialized to contain a default fragment value. The default value represents a background fragment that can be used when no fragments cover a particular subpixel sample or when all fragments that cover the particular subpixel sample are transparent. Alternatively, this default fragment value

can be stored in the graphics memory 122 where the value can be shared by multiple pixels 134. Each pixel 134 could store a special index value that pointed to the default fragment.

[0043] Other embodiments can store more than two triples in order to improve the quality of the antialiasing. Storing few triples saves memory, but can produce lesser quality antialiasing than storing many triples. For instance, it is observed that for the 8 x 8 subpixel array 202 and the sampling pattern 220 (N=8), storing three triples produces better antialiasing results than storing two triples.

[0044] Pointers 320-326 link the subpixel samples S1-S4 to the associated fragment triples 310, 312 stored in the pixel memory 314. By link, what is meant is a logical association between the subpixel samples S1-S4 and the fragment triples 310, 312. As examples, pointer 326 links subpixel S1 to fragment triple 312, while pointers 320-324 link subpixels S2-S4 to fragment triple 310.

[0045] In one embodiment, described further in connection with FIG. 6A, the linking is accomplished by storing an index value for each subpixel sample S1-S4. Accordingly, this embodiment is coined indexed sparse supersampling. In another embodiment, described in connection with FIG. 6C, the linking is accomplished by storing a coverage mask, or bit pattern, for each stored fragment value. This embodiment is hereafter referred to as an Improved A-buffer technique. Collectively, the embodiments are referred to as improved supersampling techniques.

[0046] To determine the color of the exemplary pixel 300, the graphics accelerator 108 uses one of the pixel subdivisions 200, 202, 204 and a sampling pattern 210, 220, 230 to sample the portion of the image 132 covering the pixel 300. For example in FIG. 3, the graphics accelerator 108 uses the 4 x 4 array 200 with the N = 4 sampling pattern 210 to sample pixel 300. As shown, the fragment 301 covers subpixel sample S1, and the fragment 302 covers the three subpixels samples S2-S4. A fragment covers a subpixel when the center of the subpixel sample is within an area enclosed by the fragment or, in certain cases, on an edge of the fragment.

[0047] Generally, the graphics accelerator 108 determines which fragments 301, 302 are visible at each subpixel sample S1-S4. From the perspective of a viewer of the image 132, which, for the purposes of illustrating the invention is 3D, some fragments can be closer to the viewer and in front of other fragments. The closer fragments are referred to as foreground fragments and the farther fragments, as background fragments. An opaque foreground fragment can occlude a background fragment behind that foreground fragment.

[0048] Accordingly, each fragment must pass a Z-depth test at one of the subpixel samples S1-S4, that is, the Z-value 306 of the fragment triple associated with that fragment must be smaller, i.e., closer from the perspective of the viewer, than the Z-value 306 for every other opaque fragment. If a fragment passes the Z-

depth test, then the graphics accelerator 108 stores the fragment triple associated with the visible fragment in the pixel memory 314.

[0049] When the fragment 301, for example, is determined to be visible at the subpixel sample S1 of the pixel 300, the pointer 326 is generated linking that subpixel S1 to the appropriate stored fragment triple 312. In the preferred embodiment, the pointer 326 is stored in the pixel memory 314 along with the fragment triples 310, 312 associated with the pixel 300.

[0050] Rather than storing four fragment triples in the pixel memory 314, one for each of the four subpixel samples S1-S4, which would be done using typical supersampling techniques, the exemplary embodiment in FIG. 3 stores only two fragment triples 310, 312. Accordingly, it is possible to avoid storing redundant data for the pixel 300 because only one instance of the fragment triple 310 is stored for the three subpixel samples S2-S4. By so doing, the storage requirements for fragment triples are considerably reduced.

[0051] For example, if each fragment triple 310, 312 requires nine bytes of storage, then the improved supersampling techniques use approximately eighteen bytes of memory per pixel fewer than typical supersampling methods. The improved supersampling techniques do use additional memory for storing the pointers 320-326, but this amount is small when compared to the memory saved by storing only two fragment triples 310, 312 for the four subpixel samples S1-S4.

[0052] The memory savings increase when the pixel 300 is subdivided into one of the larger subpixel arrays 202, 204. With the 8 x 8 subpixel array 202 and the sampling pattern 220 (N equals 8), the improved supersampling techniques use fifty-four fewer bytes per pixel than typical supersampling. This is because only two of eight sampled fragment triples are stored in the pixel memory 314. For the 16 x 16 subpixel array 204 and the sampling pattern 230 (N equals 16), only two of sixteen sampled fragment triples are stored in the pixel memory 314, and so 112 bytes per pixel are saved. For a display screen 130 with 1920 x 1200 pixels, such savings amount to approximately 258 Mbytes.

[0053] The displayed color of the pixel 300 depends upon which filtering function is used to combine the fragment triples associated with the four subpixel samples S1-S4. One function is simply to average the colors of the fragment triples associated with the four subpixels samples S1-S4.

[0054] FIG. 4 illustrates an exemplary case in which a third visible fragment 400 appears in the pixel 300 of FIG. 3. As indicated by an arrow 402, the third fragment 400 is linked to a new fragment triple 410. The new fragment triple 410 is different from the stored fragment triples 310, 312.

[0055] In this example, the third fragment 400 occludes a portion of fragment 302 and is visible at subpixel sample S4. The fragment 301 is still visible at the subpixel S1, as is fragment 302 at the subpixels S2-S3.

Accordingly, the subpixel sample S1 remains linked to the fragment triple 312 by the pointer 326. Subpixels S2 and S3 remain linked to the fragment triple 310 by the pointer 324 and pointer 322, respectively. To illustrate that the fragment 302 is no longer visible at the subpixel sample S4, the link 320 from the subpixel sample S4 to the fragment triple 310 is shown as broken.

[0056] When the third fragment 400 is processed by the graphics accelerator 108, the fragment triples 310, 312 are already stored in the pixel memory 314, and the fragment triple 410 is shown as not yet being stored in the pixel memory 314. Described below are various ways to handle the third fragment triple 410.

[0057] FIG. 5A shows one technique for handling the third visible fragment 400 in the pixel 300, that is, to store the corresponding fragment triple 410 in the pixel memory 314, along with the other fragment triples 310, 312. This technique presumes either that the memory 314 allocated for the predetermined number of fragment triples can accommodate the additional fragment triple 410 or that the memory 314 needed for storing the new fragment triple 410 can be dynamically allocated.

[0058] A drawback to storing additional fragment triples in the pixel memory 314 is that the amount of storage needed for the improved supersampling methods approaches or even exceeds that of typical sparse supersampling. Should a fourth fragment be visible in the pixel 300, then, in the example of the 4 x 4 subpixel array, the improved supersampling methods and sparse supersampling would each store four fragment triples. But for the larger subpixels arrays, such as the 8 x 8 array and 16 x 16 array, there is still a strong likelihood that there are fewer visible fragments in the pixel 300 than subpixel samples, and thus a corresponding savings of memory remains. Further, when pixel memory 314 is dynamically allocated beyond the predetermined number of fragment triples, in general, relatively few pixels will need dynamically allocated storage. Although improved supersampling methods might then require more storage for a given pixel 134 than typical sparse supersampling, the improved methods might use less storage for the entire image 132 overall.

#### ADAPTIVE PROCESS

[0059] Alternatively, an adaptive process can reduce the number of subpixel samples at which to sample the pixel 300 when the number of visible fragments in the pixel 300 exceeds the available storage for fragment triples, such as when the pixel memory 314 allocated for the predetermined number of fragment triples is already filled, or no pixel memory is available to dynamically allocate for the new fragment triple 410.

[0060] For example, if there is storage for only two fragment triples, but there are four different visible fragments in the pixel 300, a different fragment for each of the four subpixel samples S1-S4, then backing off to only two subpixel samples will ensure sufficient storage

11

EP 0 910 047 A2

12

for the fragments covering those two samples.

[0061] The backing off on the number of samples can be gradual. For example, if eight subpixel samples S1-S8 are used, then the process could start with eight samples, reduce to six, then four, and eventually to two, as the number of different visible fragments appear in the pixel beyond the available storage.

[0062] The process can operate independently upon each pixel. For example, the process may use all four subpixel samples S1-S4 for one pixel, and back off to only two subpixel samples S1-S2 for another pixel.

[0063] FIG. 5B illustrates still another approach for handling the third visible fragment 400 in the pixel 300, that is, to blend the corresponding fragment triple 410 with the other fragment triples 310, 312 stored in the pixel memory 314. The circled plus signs (+) in FIG. 5B illustrate the blending process.

[0064] An exemplary blending process weights the color contribution of each fragment triple 310, 312 and the new fragment triple 410 to the blended fragment triple 530, 532.

[0065] For example, the color contribution of each stored fragment triple 310, 312 is determined by multiplying the color value 304 of that fragment triple by the number of samples still covered by that fragment triple; then by dividing the result by the number of samples S1-S4 previously covered by that fragment triple before the new fragment 400 appeared. The color contribution of the new fragment triple 410 is obtained by multiplying the color value 304 of the new fragment triple 410 by the number of samples covered by the stored fragment triple, but now covered by the new fragment 400; then by dividing the result by the number of samples S1-S4 previously covered by the stored fragment triple 310, 312 before the new fragment 400 appeared.

[0066] Here, the fragment triple 310 would contribute 2/3 of its color value 304 to the blended fragment triple 530, and the new fragment triple 410 would contribute 1/3 of its color value 304. For the blended fragment triple 532, the fragment triple 312 contribute all of its color value (1/1), and the new fragment triple 410, which covers no sample points associated with the fragment triple 312, would contribute none of its color value (0/1). Then, these weighted color values 304 are added. Other color blending techniques that are known in the art can be used.

[0067] In FIG. 5B, the fragment triple 410 is blended with fragment triple 310 to produce a blended fragment triple 530, and the pointers 322, 324 linking subpixels S2-S3 to the fragment triple 310 now point to the blended fragment triple 530. Also, fragment triple 410 is blended with fragment triple 312 to produce a blended fragment triple 532, and the pointer 326 linking subpixel S1 to fragment triple 312 now points to the blended fragment triple 532. Subpixel S4 is linked to the blended fragment triple 530. Alternatively, the subpixel S4 can be linked to the other fragment triple 532.

[0068] The blended fragment triples 530, 532 are

stored in the pixel memory 314. The blended fragment triple 530 occupies the memory addresses previously occupied by the fragment triple 310. The addresses of pixel memory 314 that previously stored the fragment triple 312, now stores the blended fragment triple 532.

[0069] FIG. 5C shows an exemplary approach for accommodating the third visible fragment 400 in the pixel 300. This approach replaces one of the fragment triples 310, 312 previously stored in the pixel memory 314 with the third fragment triple 410. For example, the fragment triple 310 is replaced by the new fragment triple 410. To execute this replacement the graphics accelerator 108 would write the data of the new fragment triple 410 over the data of the previously stored fragment triple 310, in effect, discarding the data of fragment triple 310. Alternatively, memory can be deallocated for the fragment triple 310, and allocated for fragment triple 410.

[0070] In FIG. 5C, the data of the new fragment triple 410 occupies the particular addresses of pixel memory 314 that previously stored the fragment triple 310. The pointers 322, 324 point to these particular addresses of pixel memory. Where previously the pointers 322, 324 linked the subpixels S2-S3 to the fragment triple 310, these pointers 322, 324 now link the subpixels S2-S3 to the new fragment triple 410.

[0071] Techniques for selecting which fragment triples 310, 312, or 410 is discarded are described below.

### 30 Selection Schemes:

#### Z-priority

[0072] One technique for selecting the fragment triple 310, 312 to replace, called the Z-priority method, is to determine which fragment triple 310, 312 stored in the pixel memory 314 has the largest Z-depth value 306. From the perspective of a viewer, the greater the Z-depth value 306, the farther is the corresponding fragment from the viewer. For example, if the Z-depth value 306 of the fragment triple 310 is 4 and the Z-depth value 306 of the fragment triple 312 is 2, then fragment triple 310 is replaced by the new fragment triple 410. The pointers 322-324 that previously linked subpixel samples S2-S3 to the fragment triple 310 now link the subpixel samples to the fragment triple 410. In the event that more than one stored fragment triple 310, 312 has the largest Z-depth value 306, the fragment triple 310, 312 with the fewer pointers 320-326 can be replaced.

#### 50 Basic Color Difference

[0073] Another technique for selecting which fragment triple 310, 312 to replace, called the basic color difference method, involves determining which fragment triple 310, 312 stored in the pixel memory 314 has a color value 304 that is most like the color value 304 of the new fragment triple 410 i.e., produces the smallest color dif-

13

EP 0 910 047 A2

14

ference. The color value 304 of the new fragment triple 410 is compared with the color value 304 of each stored fragment triple 310, 312. Although described below using the RGB color model, this method can be applied to other color models, such as the Hue, Lightness and Saturation (HLS) and the Hue, Saturation and Value (HSV) color models.

[0074] More specifically, the basic color difference method compares the 10-bit value for the RED channel of the new fragment triple 410 with the 10-bit value for the RED channel of each stored fragment triple 310, 312. Comparisons are also made for the GREEN and BLUE channels. Values in the Alpha channels are not compared.

[0075] The absolute values of the differences between the values of the channels of the new fragment triple 410 and the values of the channels of the stored fragment triples 310, 312 are summed. Then, the sum is multiplied by the number of subpixel samples that point to the stored fragment triple 310, 312. This produces a total color difference that would result if that stored fragment triple 310, 312 were replaced by the new fragment triple 410. The fragment triple 310, 312 that produces the smaller color difference is replaced by the new fragment triple 410.

[0076] Using an overly simplified example with reference to FIG. 5C, the fragment triple 310 has a RED value of 0, GREEN value of 2, and a BLUE value of 4; the fragment triple 312 has a RED value of 2, a GREEN value of 4, and a BLUE value of 0; and the new fragment triple 410 has a RED value of 0, a GREEN value of 3, and a BLUE value of 3. Also, as shown in FIG. 5C, there are two subpixels pointing to the fragment triple 310 -- when the new fragment 400 is determined to be visible at sample point S4, the pointer 320 (see FIG. 3) from S4 to fragment triple 310 is invalidated -- and one subpixel pointing to the fragment triple 312.

[0077] The total color difference between fragment triple 310 and new fragment triple 410 is 4, e.g.,  $(|0-0| + |2-3| + |4-3|) \cdot 2$ , and the total color difference between fragment triple 312 and new fragment triple 410 is 6, e.g.,  $(|2-0| + |4-3| + |0-3|) \cdot 1$ . Thus, the fragment triple 310 is therefore replaced.

#### Color Difference and Transparent Fragments

[0078] When transparent fragments are involved in color difference, the impact of each possible replacement upon the final pixel color is compared to the ideal final pixel color that would result when the new fragment triple 410 could be stored in the pixel memory 314. That stored fragment triple 310, 312, which produces a final pixel color with the smallest color difference when compared to the ideal final pixel color, is selected for replacement. In a stack of transparent fragments, this selection tends to replace the more distant transparent fragments that are hard to see.

#### N-Squared Color Difference

[0079] In addition to comparing the new fragment triple 410 with each stored fragment triple 310, 312, as is done in the color difference method, the N-squared color difference method compares each stored fragment triple 310, 312 against each other. This method either replaces one of the stored fragment triples 310 with the new fragment triple 410, or replaces one of the stored fragment triples 310, 312 with another of the stored fragment triples 310, 312, i.e., by changing the pointers from the one stored fragment triple to that other stored fragment triple. The new fragment triple 410 is written at the addresses of pixel memory where the replaced fragment triple was previously stored. The N-squared color difference does not appear to perform significantly better than the color difference process.

#### Visual Sensitivity Color Difference Methods

[0080] Other techniques that may yield satisfactory results rely on the characteristics of the human visual system. For example, the ability of a human eye to distinguish changes in brightness may be less than the ability to perceive changes in hue. Accordingly, an exemplary visual sensitivity replacement scheme can capitalize on this characteristic by replacing the fragment triple 310, 312 that is brighter or dimmer than the new fragment triple 410 instead of the fragment triple 310, 312 that has a different hue. Such a method would prefer to replace the stored fragment triple 310, 312 with a color value 304 that differs equally, or almost equally, in each of the RGB color channels from the color value 304 of the new fragment triple 410.

[0081] Another exemplary technique can rely on the logarithmic behavior of luminance perception in humans. In general, a human eye can detect, approximately, a 2% change in the brightness of a color. Consequently, large numerical differences between high color values 304 (i.e., colors of high intensity) can be less noticeable than small numerical differences between low color values (i.e., colors of low intensity). So luminance differences are computed as ratios of color values 304, rather than as numerical differences between color values. The fragment triple 310, 312 that produces the lower luminance differences, i.e., the smaller ratio of colors, when compared to the new fragment triple 410 is replaced.

#### Z-priority Color Difference

[0082] This technique combines the Z-priority method with any of the above mentioned color difference methods to produce a replacement scheme that can perform better than any of the methods alone. The above-described color difference methods operate to replace a stored fragment triple with the new fragment triple. The Z-priority Color Difference method considers addition-



15

EP 0 910 047 A2

16

ally whether one of the stored fragment triples 310, 312 should instead replace the new fragment triple 410.

[0083] Here, the method computes color differences between the new fragment triple 410 and each stored fragment triple 310, 312 that may replace the new fragment triple 410. These color differences are computed for each of those stored fragment triples that are in front of the new fragment, i.e., lower Z-depth value, but not for those stored fragment triples that are behind the new fragment.

[0084] Accordingly, a stored fragment triple 310, 312 may be selected to replace the new fragment triple 410 when that fragment triple 310, 312 produces the smallest color difference and that fragment triple 310, 312 is associated with a fragment that is in front of the new fragment. In this case, replacement means that each subpixel sample covered by the new fragment are linked to the selected stored fragment triple 310, 312, and the new fragment triple 410 is discarded.

[0085] In general, if more than one replacement is possible, then the replacement affecting the fewer number of subpixel samples should occur. For example, if either the new fragment triple or a stored foreground fragment triple can be replaced, then the stored fragment triple replaces the new fragment triple if the stored foreground triple covers more subpixel samples than the new fragment triple.

#### Area Coverage

[0086] Another effective process selects the fragment triple that is visible at the fewest number of subpixel samples, and replaces that fragment triple with the new fragment triple 410. Afterwards, each pointer to the replaced fragment triple points to the new fragment triple.

#### Semaphore Process

[0087] The Z-priority Color Difference method allows existing foreground fragments to replace new background fragments, but does not allow existing background fragments to replace new foreground fragments. This is done to avoid losing large foreground surfaces that are made up of small foreground fragments - the large surface could be lost if the process allowed each of the small foreground fragments to be replaced by a larger background fragment. The Semaphore Process also avoids this problem.

[0088] The Semaphore Process associates a semaphore bit with each pixel. Initially, each semaphore bit is set to 0. If it is determined that replacing a new foreground fragment with an existing background fragment produces a smallest color difference, and the associated semaphore bit is 0, then the semaphore process allows the existing background fragment to replace the new foreground fragment. The associated semaphore bit is set to 1. This ensures that two such replacements

cannot occur consecutively. If the replaced new foreground fragment was part of a larger foreground surface, then the next new foreground fragment for that larger surface will replace the existing background fragment because the semaphore bit is a 1. However, it was observed that this basic semaphore process can produce some unsatisfactory artifacts.

#### Fragment Centroid Distance Methods

[0089] Such methods base color replacement on the distance between the new fragment and each possible fragment that the new fragment can replace. Accordingly, a new fragment can be extended to cover adjacent subpixel samples rather than replace stored fragments that cover distant subpixel samples. Further, it is likely that subpixel samples near the covered subpixel samples will later become covered.

[0090] Figs. 6A and 6B illustrate an exemplary logical representation of the pixel memory 314 used by the indexed sparse supersampling technique. The pixel memory 314 includes indices 600 and fragment triples 310, 312. The pixel memory 314 provides storage for a predetermined number of fragment triples. Although shown to be contiguous in the graphics memory 122, the indices 600 can be separate from each other or from the fragment triples 310, 312.

[0091] The indices 600 indicate where in the pixel memory 314 the fragment triple associated with each subpixel sample can be found. Each index 602-608 of the indices 600 is associated with one of the subpixel samples S1-S4. For example, as shown in FIG. 6B, index 602 is associated with subpixel sample S1, index 604 is associated with subpixel sample S2, and so forth. For the sampling pattern 220, there are eight indices for the eight subpixel samples S1-S8, and for sampling pattern 230, where N = 16, there are sixteen.

[0092] The value stored in each index 602-608 points to one of the fragment triples 310, 312. Accordingly, each index 602-608 links the associated subpixel sample S1-S4 to one of the fragment triples 310, 312.

[0093] When two fragment triples 310, 312 are stored in the pixel memory 314, then each index 602-608 can be represented by one data bit. The bit value stored in each index 602-608 directs the graphics accelerator 108 to the fragment triple 310, 312 associated with each subpixel sample S1-S4. In the example shown in FIG. 6B, a "1" is stored in index 602, and a "0" in each of the other indices 604-608. A zero bit value points to the first fragment triple 310 in the pixel memory 314, and a one bit value points to the second fragment triple 312.

[0094] If, alternatively, there are three fragment triples stored in the pixel memory 314, then two bits per index 602-608 are needed. Two bits per index 602-608 can accommodate as many as four stored fragment triples, three bits, as many as eight triples, and four bits, as many as sixteen.

[0095] With one bit per index 602-608, the sampling

17

EP 0 910 047 A2

18

pattern 210 (N=4) needs four bits of pixel memory 314 to implement the indices 600. The storage requirements for indices 600 of larger sampling patterns 220, 230 is also small. For example, the sampling pattern 230 (N = 16) would need 16 bits per pixel 134 to implement one bit per index. Implementing four bits per index uses 64 bits per pixel, which still provides a sizable storage savings over typical sparse supersampling techniques that store sixteen fragment triples for the sixteen subpixels samples S1-S16.

[0096] To compute a color for the pixel 300, the color value 304 of each stored fragment triple 310, 312 is multiplied by the percentage of subpixel samples linked by an index to that fragment triple. Then these weighted color values are added together to produce the pixel color.

[0097] Figs. 6C and 6D illustrate an exemplary logical representation of the pixel memory 314 used by the improved A-buffer technique. The pixel memory 314 includes coverage masks (or bit patterns) 620, 622 and stored fragment triples 310, 312. The pixel memory 314 provides storage for a predetermined number of fragment triples. Although shown to be contiguous in the graphics memory 122, the coverage masks 620, 622 can be separate from each other or from the fragment triples 310, 312. An alternative embodiment, including a third coverage mask 624 and the third stored fragment triple 410, is illustrated in Figs. 6C and 6D with dashed lines.

[0098] The coverage masks 620, 622 link the subpixel samples S1-S4 to the fragment triples 310, 312 stored in the pixel memory 314. There is one coverage mask 620, 622 associated with each stored fragment triple 310, 312. Referring to FIG. 6D, the coverage mask 620 is associated with fragment triple 310, for example, as indicated by arrow 621, and the coverage mask 622 is associated with fragment triple 312 as indicated by arrow 623. In the illustrated alternative embodiment, the coverage mask 624 is associated with the third fragment triple 410 by arrow 625.

[0099] Each coverage mask 620, 622, 624 includes one bit for each subpixel sample S1-S4. In FIG. 6D, the associations between the bits in the coverage masks and the subpixel samples S1-S4 are represented by arrows 634-640. For example, subpixel sample S1 is associated with bits 626, subpixel sample S2 with bits 628, sample S3 with bits 630 and sample S4 with bits 632.

[0100] With one bit per sample S1-S4, the sampling pattern 210 (N=4) needs four bits 626, 628, 630, 632 of pixel memory 314 to implement each coverage mask 620, 622, 624. In the shown alternative embodiment, in which three fragment triples are stored, the combined requirement for the three associated coverage masks 620, 622, 624 is twelve bits.

[0101] For the sampling pattern 220, each coverage mask 620, 622, 624 requires eight bits, one for each of the eight subpixel samples S1-S8. As for the sampling

pattern 230, which has sixteen subpixel samples S1-S16, there would be sixteen such bits in each coverage mask. Yet even with 16 bits per coverage mask, the storage savings are sizable over known sparse supersampling techniques that store sixteen fragment triples for the sixteen subpixels samples S1-S16.

[0102] The value stored in each bit of a given coverage mask indicates whether the subpixel sample associated with that bit is linked to the fragment triple associated with the given coverage mask. When a sample is linked to a fragment triple, this means that the fragment associated with that fragment triple is visible at that sample.

[0103] In the example shown in FIG. 6D, a bit pattern "0 1 1 1" is stored in the coverage mask 620. The "1" value in bits 628, 630 and 632 of the coverage mask 620 link the subpixel samples S2-S4 to the stored fragment triple 310, indicating that the fragment 302 is visible at those sample points S2-S4. Conversely, the "0" value in bit 626 of the coverage mask 620 indicates that the fragment 302 is not visible at the subpixel sample S1. Clearly, the role of each bit value can be reversed so that the "1" bit value indicates that the fragment is not visible at a given sample point, and that the "0" bit value indicates that the fragment is visible.

[0104] In the alternative embodiment shown in FIG. 6D, a third coverage mask 624 links the subpixel sample S4 to the third fragment triple 410 stored in the pixel memory 314. The association between the third coverage mask 624 and the third fragment triple 410 is noted by arrow 625.

[0105] The exemplary bit pattern stored in coverage mask 624 is "0 0 0 1", indicating that the third fragment 400 is visible at sample S4 only. Recall from FIG. 4 that new fragment 400 is linked to the fragment triple 410. For the purposes of the following illustration example, the third fragment 400 is treated as transparent. (Note that if the third fragment 400 was opaque, as described in connection with FIG. 4, then the bit pattern in the coverage mask 620 would change from "0 1 1 1" to "0 1 1 0" to indicate that the fragment 400 occluded the fragment 302 at the subpixel sample S4.)

[0106] Because the third fragment 400 is transparent, two coverage masks 620, 624 link the subpixel sample S4 to two stored fragment triples 310, 410. The "1" bit values in bit positions 632 of the coverage mask 620 and 624 indicate that both fragments 302 and 400 are visible at the subpixel sample S4. Generally, any subpixel sample S1-S4 can be linked to multiple stored fragment triples, where one fragment is opaque and each other fragment is transparent. In fact, all of stored fragment values can be transparent when a default background fragment value is used.

[0107] Accordingly, the Improved A-buffer technique described herein can support order-independent transparency, i.e., the system 100 does not need to partition primitives of the image 132 so as to present transparent fragments to the graphics accelerator 108 after present-

19

EP 0 910 047 A2

20

ing all opaque fragments, nor does the system 100 need to sort the transparent primitives so as to present transparent fragments in Z-depth order.

[0108] To compute the color of the pixel 300, a color is first computed for each subpixel sample S1-S4, and then the computed colors are combined. Where a subpixel sample is linked to one opaque fragment only, such as sample S1, the color for that subpixel sample S1 is the color of the associated stored fragment value 312.

[0109] Where a subpixel sample, such as sample S4, is linked to two stored fragment triples 310, 410, one transparent 400 and the other opaque 302, the color for the subpixel sample S4 is the sum of the color contributions of those two fragment triples 310, 410. The color contribution of the transparent fragment 400 is the opacity of that fragment 400, as indicated by the value stored in the Alpha channel, multiplied by the color of that fragment 400. The contribution C of the opaque fragment 302 is the color of that fragment f(c) 302 multiplied by 1 minus the opacity of the transparent fragment f(o). 400, e.g.,  $C = f(c) \times (1 - f(o))$ .

[0110] The exemplary embodiments shown in Figs. 6A-6D can achieve satisfactory antialiasing results by storing two fragment triples for four subpixel samples. Eight subpixel samples with two stored fragment triples usually looks better than four subpixel samples with two fragment triples, but can look worse when one of the additional four subpixel samples requires replacing one of the stored triples with a third fragment triple, and that third fragment triple appears in the pixel memory last. Thus, allocating storage for a third fragment triple can make a marked improvement for eight subpixel samples over storing two fragment triples. Clearly, the antialiasing results can be made to approach the results of typical sparse supersampling as more fragment triples are stored, but each additional triple erodes the memory savings provided by the improved supersampling techniques.

[0111] Fig. 7 illustrates a flow diagram 700 describing the process of generating an image 132 using the present invention. In the early stages of processing the image 132, the image 132 is partitioned into fragments. When processing each new fragment, the graphics accelerator 108 determines whether the new fragment is visible at any of the subpixel samples S1-S4 covered by the new fragment. The graphics accelerator 108 compares the Z-depth value of the new fragment with the Z-depth value of the stored fragment associated with each covered subpixel sample S1-S4 (step 702).

[0112] If the new fragment has a smaller Z-depth value than the Z-depth value of a stored fragment for any covered subpixel sample S1-S4, then the new fragment is in front of that stored fragment and, consequently, is visible. An exception, however, is when the new fragment has an Alpha value of 0.0. In this instance the new fragment is completely transparent. The graphics accelerator 108 does not need to store the fragment

value of the new fragment because the new fragment is, in effect, invisible.

[0113] If instead the new fragment has a larger Z-depth value than the Z-depths values for all of the covered subpixel samples S1-S4, then the new fragment is behind one or more stored fragments and may be invisible. If the new fragment is behind opaque foreground fragments, then the new fragment is invisible, and the processing of the new fragment for the pixel 134 is complete. If, however, the new fragment is immediately behind a transparent foreground fragment, then the new fragment can still be seen.

[0114] When the new fragment is visible at one of the covered subpixel samples, then the graphics accelerator 108 invalidates the link between each covered sample and a stored fragment, if the new fragment obscures the stored fragment for that covered subpixel sample. For the indexed sparse supersampling technique, the graphics accelerator 108 maintains control bits for keeping track of the validity of each index and invalidates each index linking a covered subpixel sample to an obscured fragment. The control bits may direct the graphics accelerator 108 to use the default background color if no fragments cover a subpixel sample. For the improved A-buffer technique, the bits in the coverage mask associated with each covered subpixel sample are unchanged when the new fragment is transparent and are set to "0" when the new fragment is opaque.

[0115] Then, in step 708, the number of links pointing to each fragment triple is counted. For the indexed sparse supersampling technique, step 708 counts the number of indices linked to that stored fragment triple. For the improved A-buffer technique, step 708 counts the number of bits in each coverage mask that have a "1" bit value.

[0116] A fragment triple is free and available when there are no links pointing to the fragment triple. In this case a new fragment triple associated with the new fragment can replace that free fragment triple.

[0117] If step 710 determines that a fragment triple is free, then the new color associated with the new fragment is stored in the freed fragment triple (step 712). In step 714, the links of the subpixel samples covered by the new fragment are set to the new fragment triple.

[0118] If step 710 determines that no fragment triples are free, then a replacement scheme, such as the color difference technique described above, selects one of the stored fragment triples for replacement (step 716). Replacement means changing the color, Z-depth, and stencil values stored in the selected fragment triple to the color, Z-depth, and stencil values of the new fragment triple.

[0119] In step 718, the new color is written to the selected fragment triple. The links originally pointing to the selected fragment triple are still pointing to that fragment triple. Because the selected fragment triple now contains a value representing the new color, the subpixel samples associated with such links are thereby

21

EP 0 910 047 A2

22

associated with the new color. Those links corresponding to the subpixel samples covered by the new fragment are set to point to the new color (step 714).

[0120] In step 720, the pixel color is computed from the subpixel samples as described above in connection with Figs. 6A-6D. The links associated with the subpixel samples S1-S4 point to the stored colors that are used to produce the color of the pixel 134. Accordingly, the pixel color can change as each new fragment appears in the pixel 134.

[0121] When, in step 722, the graphics accelerator 108 is through processing all fragments, then the pixels are ready for display (step 724).

[0122] In FIG. 7, an alternative process for generating an image computes the color of the pixel in step 720', illustrated as dashed lines, before determining whether there are any fragment triples available in which to store the new fragment value associated with the new fragment. The existing colors stored in the fragment triples and the new color value combine to produce the pixel color. The effect is to compute the color as though an additional triple was available.

[0123] After computing the pixel color in step 720', the alternative process may then replace an existing stored color with the new fragment triple as described above in connection with the steps 716-718. If each fragment processed after this new fragment does not lead to a new computation of the pixel color, then no color data is lost despite the replacement.

[0124] It is to be understood that the above-described embodiments are simply illustrative of the principles of the invention. Various other modifications and changes may be made by those skilled in the art which will embody the principles of the invention and fall within the scope thereof.

## Claims

1. A computerized method for rendering an image defined by pixels, comprising the steps of:

selecting subpixel positions in a pixel as sample points;  
storing a fragment value associated with a fragment of the image, the fragment covering at least one of the sample points; and  
linking each covered sample point to the stored fragment value to enable the generation of a color of the pixel using the fragment value, the generated color improving the perceived quality of the image by reducing aliasing artifacts.

2. The method of claim 1, wherein the step of linking includes the steps of:

associating an index with each covered sample point; and  
storing a value in each index that points to the

stored fragment value.

3. The method of claim 1, wherein the step of linking includes the steps of:

associating a bit pattern with the stored fragment value, the bit pattern having a sequence of bits, each bit in the sequence being associated with one of the sample points; and  
storing a value in each bit, the value in each bit indicating whether or not the sample point associated with that bit is covered by the fragment.

4. The method of claim 1 further comprising the step of:

partitioning the pixel into an array of subpixel positions, the array having N rows and N columns where N is an integer greater than one; and wherein each selected sample point is in a different row and in a different column than every other sample point in that pixel.

5. The method of claim 1, wherein there are fewer sample points than subpixel positions in the pixel

6. The method of claim 1 further comprising the step of:

allocating memory to the pixel for storing a predetermined number of fragment values, the predetermined number being less than the number of sample points in the pixel, further comprising the step of:

determining whether memory is available for storing the fragment value by counting links to each previously stored fragment value.

7. The method of claim 1 further comprising the steps of:

processing a new fragment having a new fragment value, the new fragment being visible at one or more of the sample points of the pixel; and  
replacing the stored fragment value with the new fragment value, further comprising the step of:

computing a color for the pixel using the new fragment value and each stored fragment value before replacing the stored fragment value with the new fragment value.

23

EP 0 910 047 A2

24

8. The method of claim 1 further comprising the step of:

reducing a number of sample points for the pixel when storing a fragment value associated with a subsequent fragment would exceed an amount of memory available for the pixel.

9. The method of claim 1 further comprising the steps of:

selecting subpixel positions in a second pixel as sample points; and  
linking one of the sample points of the second pixel to the stored fragment value.

10. A method for reducing aliasing artifacts in an image defined by pixels comprising the steps of:

selecting subpixel positions in a pixel as sample points;  
determining that a new fragment of the image is visible at one or more of the sample points of the pixel, the new fragment having a new fragment value;  
allocating addresses of memory for the pixel for storing a predetermined number of fragment values;  
replacing a selected one of the stored fragment values with the new fragment value if the predetermined number of fragment values are already stored in the allocated memory when the new fragment is determined to be visible, otherwise, storing the new fragment value;  
linking each sample point of the pixel at which the new fragment is visible to the stored new fragment value to enable the generation of a color of the pixel using the fragment value, the generated color improving the perceived quality of the image by reducing aliasing artifacts.

11. A memory for storing data used to compose a color for a pixel that is partitioned into subpixels, comprising:

a plurality of addresses, each address storing a fragment value, each stored fragment value including a color value and being associated with a fragment of an image that is visible at one of the subpixels; and  
another address storing indices, each index being associated with one subpixel and including an index value that points to one of the plurality addresses storing a particular fragment value to link the subpixel associated with that index with the color value of the particular fragment value.


12. An apparatus for rendering an image defined by pixels comprising:

means for selecting subpixel positions in a pixel as sample points;  
means for storing a fragment value associated with a fragment of the image, the fragment covering one or more of the sample points in the pixel; and  
means for linking each covered sample point to the stored fragment value to enable the generation of a color of the pixel using the fragment value, the generated color improving the perceived quality of the image by reducing aliasing artifacts.

13. An apparatus for rendering an image defined by pixels comprising:

a graphics device selecting subpixel positions in a pixel as sample points; and  
memory, coupled to the graphics device, storing a fragment value associated with a fragment of the image, the fragment covering one or more of the sample points in the pixel, the graphics device linking each covered sample point of the pixel to the stored fragment value to enable the generation of a color of the pixel using the fragment value, the generated color improving the perceived quality of the image by reducing aliasing artifacts.

EP 0 910 047 A2



EP 0 910 047 A2




FIG. 2A




FIG. 2B




FIG. 2C

EP 0 910 047 A2




FIG. 3



EP 0 910 047 A2




FIG. 4

**FIG. 5A**

EP 0 910 047 A2




FIG. 5B

EP 0 910 047 A2




FIG. 5C

EP 0 910 047 A2






FIG. 6A

EP 0 910 047 A2




EP 0 910 047 A2



EP 0 910 047 A2

FIG. 6D





EP 0 910 047 A2

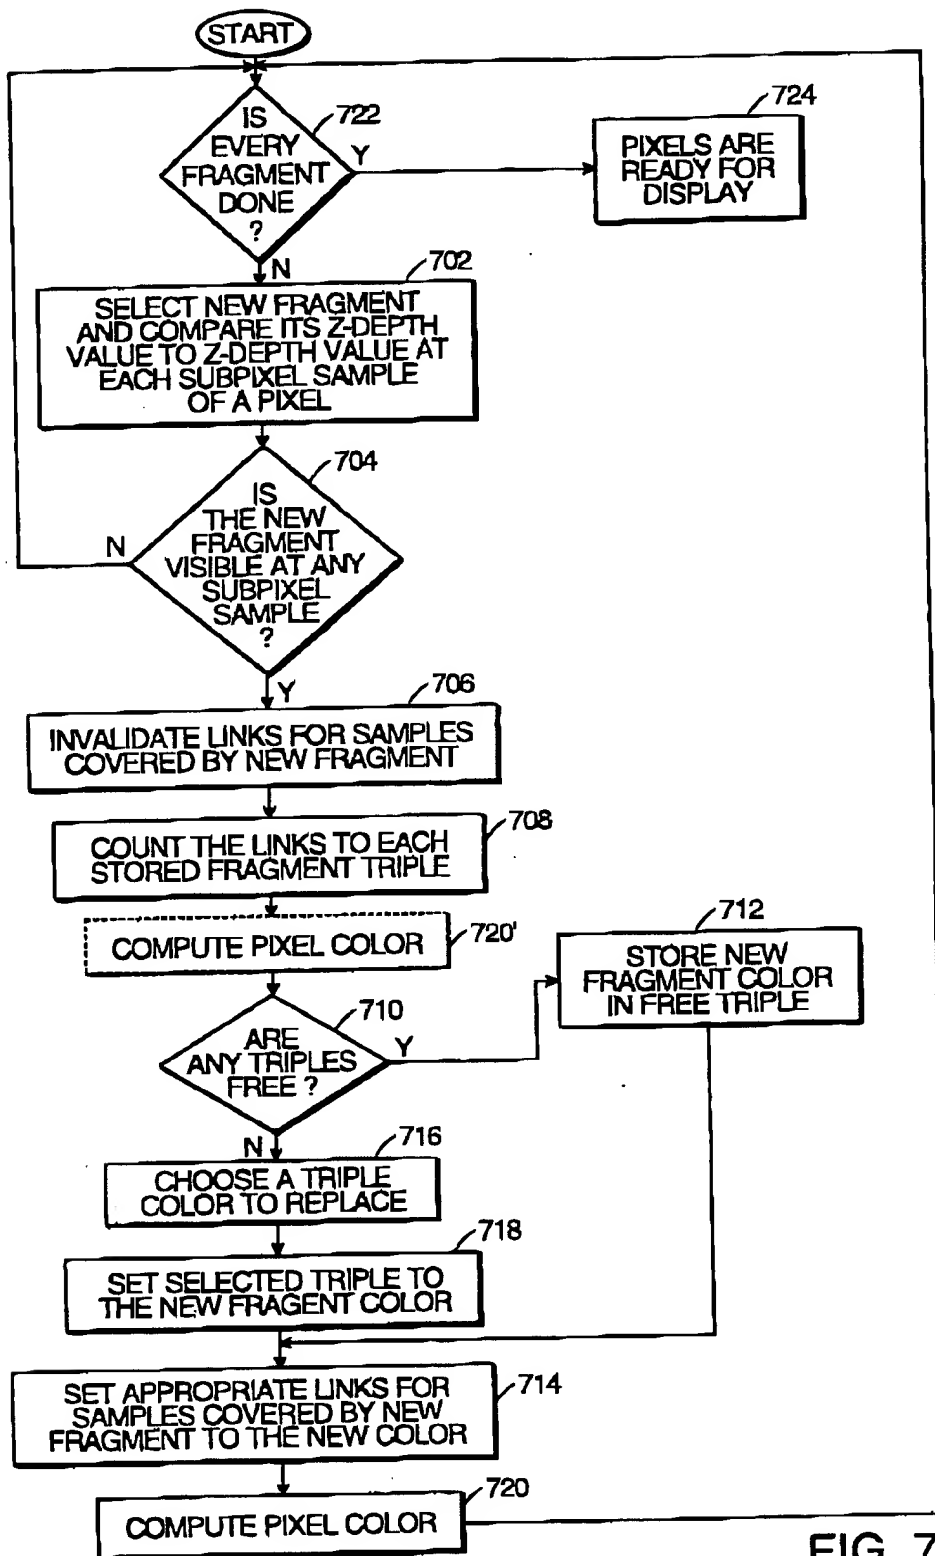


FIG. 7

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 910 047 A3

(12)

## EUROPEAN PATENT APPLICATION

(88) Date of publication A3:  
07.02.2001 Bulletin 2001/06

(51) Int. Cl.<sup>7</sup>: G06T 15/10, G06T 15/50

(43) Date of publication A2:  
21.04.1999 Bulletin 1999/16

(21) Application number: 98307790.0

(22) Date of filing: 25.09.1998

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventors:

- Jouppl, Norman, P  
Palo Alto, California 94306 (US)
- McCormack, Joel, M  
Boulder, Colorado 80302 (US)

(30) Priority: 15.10.1997 US 958129

(74) Representative:

Charig, Raymond Julian et al  
Eric Potter Clarkson,  
Park View House,  
58 The Ropewalk  
Nottingham NG1 5DD (GB)

(71) Applicant:  
Compaq Computer Corporation  
Houston Texas 77070 (US)

(54) Full-scene antialiasing using improved supersampling techniques

(57) A method and an apparatus reduces aliasing artifacts in images defined by pixels. A pixel is partitioned into subpixel locations from which sample points are selected. A fragment of the image is determined to be visible at least one of the sample points. A fragment value associated with that fragment is stored. Each

sample point at which the fragment is visible is linked to the stored fragment value. A color of the pixel is computed from the stored fragment values to reduce the aliasing artifacts in the image.




FIG. 1

EP 0 910 047 A3

EP 0 910 047 A3

European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number

EP 98 30 7790

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl. 6)
P, X	WO 97 41536 A (EVANS & SUTHERLAND COMPUTER CO.; COSMAN MICHAEL A (US); DREWS MICHA) 6 November 1997 (1997-11-06) * page 2, line 31 - page 3, line 27 * * page 5, line 29 - line 35 * * page 6, line 33 - page 7, line 32 *	1-7, 9-13	G06T15/10 G06T15/50
A	EP 0 463 700 A (PHILIPS ELECTRONICS UK LTD; PHILIPS NV (NL)) 2 January 1992 (1992-01-02) * column 2, line 42 - column 3, line 33 *	1, 4, 5, 7, 10, 13	
			TECHNICAL FIELDS SEARCHED (Int. Cl. 6)
			G06T
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 19 December 2000	Examiner Burgaud, C
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document			

EPO FORM 1503 (03/02) (P/AnC/1)

EP 0 910 047 A3

**ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.**

EP 98 30 7790

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

19-12-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9741536 A	06-11-1997	US 5818456 A	06-10-1998
		AU 2824697 A	19-11-1997
		CA 2253297 A	06-11-1997
		DE 19781737 T	12-05-1999
		GB 2328596 A	24-02-1999
		US 5943060 A	24-08-1999
EP 0463700 A	02-01-1992	GB 2245805 A	08-01-1992
		GB 2245806 A	08-01-1992
		DE 69127516 D	09-10-1997
		DE 69127516 T	26-02-1998
		JP 4233086 A	21-08-1992
		KR 239969 B	15-01-2000
		US 6088036 A	11-07-2000
		DE 69122557 D	14-11-1996
		DE 69122557 T	24-04-1997
		EP 0464907 A	08-01-1992
		JP 4233672 A	21-08-1992
		US 5394516 A	28-02-1995

EPO FORM P0198

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82